

Frequency-Domain Synthesis of Advanced Pulse-Compression Filters

Alan C. O'Connor

MIT Lincoln Laboratory, Lexington, MA, USA

Email: alan.oconnor@ll.mit.edu

Abstract—Radars require modified signal processing to be able to employ waveforms other than linear frequency modulated chirps. This paper describes a computationally-efficient method for designing pulse-compression filters that improve performance. These pulse-compression filters have been previously described in terms of a time-domain representation—finding the finite impulse response filters required solving a large system of linear equations. The approach described here translates the filter design specifications into the frequency domain. The frequency-domain representations lead to algorithms with an order of magnitude lower complexity and that are more readily parallelized. The design approaches described here are an important step to using these advanced pulse-compression filters in situations where they must be computed in real-time.

I. INTRODUCTION

Coherent pulse trains of linear frequency-modulated (LFM) waveforms combined with matched-filter pulse compression, or mismatch filtering via tapering, are the workhorses of modern surveillance radars [1]–[3]. This combination gives excellent performance for both synthetic aperture radar (SAR) and moving target indication (MTI) radars, but offers little flexibility.

Recent years have seen growing interest in waveform diversity generally, and pulse-to-pulse waveform agility more specifically. For example, pulse-to-pulse agile waveforms have been proposed as a tool to mitigate range ambiguities [4], [5], encode communication data in the radar transmission [6], or avoid interference to other users in a dynamic spectrum environment [7]. Additional potential benefits have been discussed in [8], [9], and [10].

Unfortunately, moving away from a train of identical LFM pulses can cause performance problems for a radar. The specifics are different for SAR and MTI, but in both cases, advanced pulse-compression filters can, to a degree, mitigate the loss in performance.

In MTI radars, returns from multiple pulses are combined coherently to cancel echoes from stationary ground clutter and enable detection of only moving targets. The main problem with using conventional pulse compression with pulse-to-pulse

agile waveforms in MTI radar is that matched filters will cause range-sidelobe modulation (RSM) of the clutter energy. Because the range sidelobes of the impulse response (IPR) differ on each pulse, clutter energy that leaks into these range sidelobes cannot be canceled by the subsequent processing, and prevents detection of weak targets. This clutter leakage problem can be corrected by designing pulse-compression filters with matched IPRs. This approach was given the moniker “joint least-squares” (JLS) filters in [6], and were studied further in [11] and [12].

For SAR, the higher IPR sidelobes of non-LFM waveforms cause a reduction in image contrast that is quantified by the multiplicative noise ratio (MNR). Better MNR can be obtained by using a pulse-compression filter that yields a low integrated sidelobe ratio (ISLR). “Mismatch filters” (MMF) are pulse-compression filters that minimize ISLR; they have been described in [13]–[15].

A. Contribution of this Paper

Prior presentations of MMF and JLS filters used a time-domain representation for the filter design objectives and algorithms. This paper describes frequency-domain synthesis of MMF and JLS pulse-compression filters, which is much simpler computationally, and could open the door to greater application of pulse-to-pulse waveform agility.

B. Radar Processing Background

Both SAR and MTI radars typically employ bursts of pulses called a coherent processing interval (CPI). Suppose a radar illuminates a scene with a CPI of M pulses. In the case of pulse-to-pulse agility, a different waveform s_m , $m \in \{1 \dots M\}$ is used for each pulse. The received signal is the waveform convolved with the scene response plus complex Gaussian noise added by the receiver electronics. Pulse compression is applied to the returns from each pulse. Finally, the data is coherently combined across the CPI using processing that depends on whether the radar is performing SAR or MTI [2].

The filters described in the paper are the set of FIR filters h_m used to compress returns from each pulse. The pulse-compression filters should coherently combine signal correlated with the radar waveform, while not amplifying the receiver noise excessively. For SAR, the specific objective is to design a set of FIR filters h_m , which when convolved with the corresponding waveform s_m produce an IPR with low

integrated sidelobe ratio (ISLR); this is achieved by mismatch filters. For MTI radar, the objective is to design filters that produce matched IPRs across a set of waveforms, so that all the clutter energy can be canceled; this objective leads to joint least-squares filters.

C. Example

To motivate the frequency-domain synthesis approach, and give some intuition about what the MMF and JLS filters are trying to accomplish, we start with an example. In Fig. 1, we show the waveform power spectral densities (PSD) for a set of two waveforms, each a random quadrature phase shift keyed (QPSK) signal with 25 chips. The MMF and JLS filters were selected to be three times the length of the waveforms, so have 75 coefficients. The PSDs of the filters are shown in the second row of plots; MMF in the left column, and JLS in the right column. The third row of plots are the waveform-filter cross-spectral densities (CSD), which is the fast Fourier transform (FFT) of the time-domain impulse response. The impulse responses for the MMF and JLS filters are shown in Fig. 2. Note that, in the frequency-domain, the effect of MMF is to cancel out the ripples in the waveform PSD, so that the waveform-filter CSD is quite flat. A flat CSD corresponds to an impulse response with low sidelobes. In contrast, the effect of the JLS filters is to cancel the ripples relative to a common, optimal magnitude response. Because this CSD is optimized for match across the CPI rather than flatness, the resulting IPRs are matched, but do not have low sidelobes.

II. TIME DOMAIN FILTER SYNTHESIS

The time-domain filter synthesis for MMF and JLS filters, and proofs of optimality of the filters for their corresponding objective functions are given in [11]-[12]. The derivations will not be repeated here.

We will merely mention that time-domain synthesis for both kinds of filters involves solving very large systems of equations that describe the objective functions in terms of time-domain convolution of the filters and waveforms. The superiority of frequency-domain filter synthesis approaches described here are based on the idea that convolution in the time-domain is equivalent to point-wise multiplication in the frequency-domain. Computations that are coupled together in the time-domain synthesis become uncoupled in the frequency domain, leading to algorithms that are much simpler.

As described in [12], the linear system of equations describing the JLS or MMF filters in the time-domain can also be solved via gradient descent. Because the system of equations involves convolutions, the matrix multiplications required by the gradient solver can be replaced by fast convolutions. As we will show in Section IV, this gradient descent approach is much faster than the direct solution of the system of equations. Nevertheless, frequency-domain synthesis is even faster.

III. FREQUENCY-DOMAIN FILTER SYNTHESIS

The basic idea is that the target frequency-domain content of the filters can be determined point-wise. The point-wise value is determined by dividing the target value of

the waveform-filter cross-spectral density by the waveform frequency-domain content. The difficulty with this approach is that the point-wise optimal frequency-domain content does not correspond to a time-domain FIR filter of the desired length. Simply truncating the synthesized time-domain filter taps leads to poor performance. Instead, we follow the classic “window method” filter design approach [16]. The window method is typically used to design filters with much simpler frequency responses, but since the waveform spectra are oversampled (by at least a factor of κ), the frequency-domain target is smooth enough that the time-domain filter taps are well behaved.

We next derive the target frequency-domain content for MMF and JLS filters in terms of optimization problems, and later describe the window functions that are used to obtain time-limited filters.

A. Notation

We use $(\cdot)^*$ to indicate complex conjugate, $(\cdot)^T$ for transpose, and $(\cdot)'$ for Hermitian conjugate. $\Re(\cdot)$ will indicate the real part of a complex number; $\Im(\cdot)$ will indicate discrete Fourier transform of a sequence. Let $\mathbf{I}_M$ stand for the $M \times M$ identity matrix, and $\text{diag}(\mathbf{X})$ for the matrix with the elements of vector $\mathbf{X}$ along the diagonal.

Let x_m be the m th sampled, complex-baseband waveform. The sampled waveforms have length N_c . The filters sought have length $\tilde{N}_c = \kappa N_c$, where $\kappa \geq 1$.

The clutter-to-noise ratio (CNR) from clutter in one range-bin, before Doppler processing, is σ_c^2/σ_n^2 .

B. Matched-Sidelobe JLS Filters

The JLS objective of matching IPRs across a set of M pulses can be framed as matching of filter-waveform CSDs. Let $X_m = \Im(x_m)$ be the FFT of the m th waveform of the CPI, and let $H_m^{\text{jls}} = \Im(h_m^{\text{jls}})$ be the FFT of the corresponding filter that we want to design. The FIR filter and waveform should be zero-padded, so that they have a convenient length for computing FFTs; let this FFT length be N_f .

The objective of matching IPRs is captured by the following constrained minimization problem:

$$\Gamma_{\text{jls}} = \frac{1}{2M(M-1)} \frac{\sigma_c^2}{\sigma_n^2} \sum_{m=1}^M \sum_{n=1}^M \sum_{k=1}^{N_f} |H_m[k]X_m^*[k] - H_n[k]X_n^*[k]|^2 + \frac{1}{M} \sum_{m=1}^M \|H_m\|^2, \quad (1)$$

$$\{H_1^{\text{jls}}, \dots, H_M^{\text{jls}}\} \leftarrow \arg \min \Gamma_{\text{jls}} \quad (2)$$

$$\text{subject to } \sum_{k=1}^{N_f} X_m[k] \cdot H_m[k] = N_f, \quad \forall m = 1 \dots M \quad (3)$$

The first term of (1) penalizes mismatch between the CSDs, and the second term penalizes increased filter noise gain. The M constraints (3) ensure that the zero-delay sample of each time-domain impulse response is unity.

We now rearrange the summations in the objective (1) by stacking the coefficients involved at a particular frequency k so that $\mathbf{X}[k] = (X_1[k], \dots, X_M[k])^T$ is the vector of waveform




Fig. 1. Filter and waveform power spectral densities and corresponding impulse responses, showing the frequency domain content of two waveforms (a), the corresponding MMF and JLS filters (b,e), and the waveform-filter cross-spectral densities for each (c,f). Subplot (a) is duplicated for ease of comparison. The impulse responses for these waveforms and filters are shown in Fig. 2.




Fig. 2. Impulse responses produced by MMF and JLS filters for the waveforms in Fig. 1.

coefficients and $\mathbf{H}[k]$ is the corresponding vector of JLS filter coefficients at frequency k .

Using the stacked notation, the objective function can be rewritten as a sum of quadratic forms, one at each frequency:

$$\Gamma_{\text{jls}} = \frac{1}{2} \sum_{k=1}^{N_f} \mathbf{H}'[k] \Sigma_{\text{S}}[k] \mathbf{H}[k] \quad (4)$$

where $\Sigma_{\text{S}}[k]$ is an $M \times M$ matrix equal to:

$$\eta \cdot \left(M \cdot \text{diag}(\mathbf{X}[k]) \text{diag}(\mathbf{X}^*[k]) - \mathbf{X}[k] \mathbf{X}'[k] \right) + \epsilon \mathbf{I}_M \quad (5)$$

and where $\eta = 2(\sigma_c^2/\sigma_n^2)/(M(M-1))$ and $\epsilon = 2/M$.

Finally, we incorporate Lagrange multiplier terms to enforce the normalization constraints in (3). Let $\mathbf{1}$ be a vector of M ones, and Λ a vector of the M Lagrange multipliers. In this notation, the unconstrained objective is:

$$\tilde{\Gamma}_{\text{jls}} = \frac{1}{2} \sum_{k=1}^{N_f} \mathbf{H}'[k] \Sigma_{\text{S}}[k] \mathbf{H}[k] + \Re \left\{ \Lambda' \left(N_f \cdot \mathbf{1} - \sum_{k=1}^{N_f} \text{diag}(\mathbf{X}[k]) \cdot \mathbf{H}[k] \right) \right\} \quad (6)$$

By equating the Wirtinger gradient [17] of $\tilde{\Gamma}_{\text{jls}}$ with respect to $\mathbf{H}[k]$ to zero, the minimum for frequency index k can be shown to occur at:

$$\mathbf{H}_{\text{jls}}[k] = \Sigma_{\text{S}}^{-1}[k] \cdot \text{diag}(\mathbf{X}^*[k]) \cdot \Lambda \quad (7)$$

Substituting this into the constraint (3), the normalization constants Λ can be found to be:

$$\Lambda = N_f \cdot \left(\sum_{k=1}^{N_f} \text{diag}(\mathbf{X}[k]) \cdot \Sigma_{\text{S}}^{-1}[k] \cdot \text{diag}(\mathbf{X}^*[k]) \right)^{-1} \cdot \mathbf{1} \quad (8)$$

C. Low-Sidelobe Mismatch Filters

Mismatch filters (MMF) result when we seek pulse-compression filters that minimize the integrated sidelobe ratio (ISLR) of the impulse responses.

As above, let X_m stand for the FFT of the m th waveform, and H_m for the FFT of the coefficients of the mismatch filter that we want to design.

Given M waveforms, we seek a set of corresponding pulse-compression filters that solve the following constrained minimization:

$$\Gamma_{\text{mmf}} = \frac{1}{M} \cdot \frac{\sigma_c^2}{\sigma_n^2} \cdot \sum_{m=1}^M \sum_{k=1}^{N_f} |H_m[k] X_m^*[k]|^2 + \frac{1}{M} \cdot \sum_{m=1}^M \|H_m\|^2, \quad (9)$$

$$\{h_1^{\text{mmf}}, \dots, h_M^{\text{mmf}}\} \leftarrow \arg \min \Gamma_{\text{mmf}} \quad (10)$$

$$\text{subject to } \sum_{k=1}^{N_f} X_m[k] \cdot H_m[k] = N_f, \forall m = 1 \dots M \quad (11)$$

The objective function balances the need for low ISLR (first term) with a that for low noise gain (second term), because both will increase the SAR image background level.

The objective function (9) can be rewritten in the stacked, point-wise frequency notation as:

$$\Gamma_{\text{mmf}} = \frac{1}{2} \sum_{k=1}^{N_f} \mathbf{H}'[k] \Sigma_{\text{D}}[k] \mathbf{H}[k] \quad (12)$$

where $\Sigma_{\text{D}}[k]$ is an $M \times M$ matrix equal to:

$$\Sigma_{\text{D}} = \tilde{\eta} \cdot \left(\text{diag}(\mathbf{X}[k]) \cdot \text{diag}(\mathbf{X}^*[k]) \right) + \epsilon \mathbf{I}_M, \quad (13)$$

and where $\tilde{\eta} = 2(\sigma_c^2/\sigma_n^2)/M$ and $\epsilon = 2/M$. Finally, as for the JLS filters, Lagrange multipliers are added to the objective function to enforce the normalization constraints (11):

$$\tilde{\Gamma}_{\text{mmf}} = \frac{1}{2} \sum_{k=1}^{N_f} \mathbf{H}'[k] \Sigma_{\text{D}}[k] \mathbf{H}[k] + \Re \left\{ \Lambda' \left(N_f \cdot \mathbf{1} - \sum_{k=1}^{N_f} \text{diag}(\mathbf{X}[k]) \cdot \mathbf{H}[k] \right) \right\} \quad (14)$$

The MMF are given by the minimizer of (14), which is:

$$\mathbf{H}_{\text{mmf}}[k] = \Sigma_{\text{D}}^{-1}[k] \cdot \text{diag}(\mathbf{X}^*[k]) \cdot \Lambda \quad (15)$$

Since Σ_D is diagonal, the filters can be computed independently. The k th Fourier component of the m th MMF filter is related to the k th Fourier component of the m th waveform by:

$$H_m^{\text{mmf}}[k] = \frac{X_m^*[k] \cdot \lambda_m}{\tilde{\eta} \cdot |X_m[k]|^2 + \epsilon} \quad (16)$$

The Lagrange multipliers $\Lambda = (\lambda_1 \dots \lambda_M)^T$ enforce normalization of each filter impulse response, and are independent.

D. Windowing Functions

The time-domain filter taps for each filter h_m are then obtained via inverse FFT and are multiplied by a window function:

$$h_m = w \odot \mathfrak{F}^{-1}(H_m) \quad (17)$$

Recall that the FFT length N_f was chosen to be longer than the desired number of taps. Using a window function when truncating the taps avoids introducing ripples in the spectrum [16]. To obtain filters more similar to those obtained via the time-domain method, only those coefficients that extend beyond the sample support of the waveform should be shaped by the window function—the central N_c coefficients should be left unmodified. The remaining $(\kappa - 1)N_c$ coefficients can be shaped by a splitting a conventional window function, such as a Kaiser window. Because the central taps are left unmodified, the normalization conditions (3) and (11) are still satisfied after windowing.

E. Frequency-domain Tapering

Extra care is required in constructing MMF when the waveforms are oversampled (i.e. faster than Nyquist sampling). In the time-domain synthesis, good sidelobe suppression is obtained only if certain rows of the waveform correlation matrices are zeroed out, so as not to perturb the IPR mainlobe [18]. In the frequency-domain construction of MMF, this problem is mitigated by applying a frequency-domain taper. We find that if the waveform is oversampled by a factor Θ , tapering with the function

$$\Omega(f) = \left| \frac{\sin(\pi f \Theta)}{\pi f \Theta} \right|, \quad f \in [-0.5, 0.5) \quad (18)$$

gives a similar response to the time-domain synthesis.

For JLS filters, it is not necessary to zero out any rows of the waveform correlation matrices when synthesizing in the time domain [12], nor is it necessary to apply a taper when using the frequency-domain approach.

F. Selection of Filter Length Multiple κ

Larger values of κ yield better theoretical filter performance. However, in a real employment where the extent of data recorded between successive pulses is finite, larger κ increases the range-extent over which edge-effects degrade the actual achieved sidelobe performance. Longer filters also have worse Doppler tolerance [12], [14].

IV. RESULTS

A. Comparison of Time- versus Frequency-Domain Synthesis

With appropriate selection of window function, frequency-domain synthesis produces taps with performance quite similar to the time-domain approach. In Fig. 3 we show how the ISLR of the mismatch filter designed in the frequency domain depends on the window parameter. For a Kaiser window parameter near 1.5, frequency-domain MMF achieves an ISLR within 1 dB of the filter synthesized in the time domain. The results shown are the average ISLR for 50 random QPSK sequences, each of length 300. The approximation is better for larger κ . By comparison, pulse compression with a matched filter leads to $\text{ISLR} \approx 0$ dB for random QPSK waveforms.




Fig. 3. Integrated sidelobe ratio performance of frequency-domain synthesized MMF versus Kaiser window parameter. Results are shown for $\kappa = 2$ (squares), $\kappa = 3$ (circles), and $\kappa = 5$ (triangles). ISLR for matched filter and time-domain synthesized MMF are shown for comparison.

The taps obtained by the two methods are quite similar—the correlation coefficient for the filter tap sequences obtained by the two methods is typically above 0.99. Example MMFs are shown in Fig. 4.




Fig. 4. Mismatch filter tap amplitudes synthesized via time-domain and frequency-domain methods. Example waveform was 30-sample QPSK sequence. Filter is three times longer than waveform ($\kappa = 3$). Filter taps generated by the two methods have correlation coefficient of 0.994.




Fig. 5. Scaling of complexity for three JLS synthesis methods versus filter length κN_c , for fixed number of waveforms $M = 7$.

B. Comparison of Computational Complexity

We now examine how the time to synthesize JLS filters scales with filter length κN_c and number of waveforms M . For the time-domain synthesis, we considered both the direct solution of the linear system involving the waveform correlations, as well as a gradient descent method. We recorded computation times for Matlab implementations of each algorithm; results are shown in Figs. 5–6. We now examine how complexity scales for the step in each algorithm that was the primary driver of computation time.

For the direct time-domain approach, complexity is driven by solving a linear system of size $M\kappa N_c$, which has complexity:

$$C_{\text{TD, direct}} = \mathcal{O}(M^3 \kappa^3 N_c^3) \quad (19)$$

A more efficient time-domain approach uses gradient descent and furthermore takes advantage of the block structure inherent in the problem to replace matrix multiplications in the gradient computation with fast FFT-based convolutions [12]. Each matrix multiply involves M^2 blocks, and the system of equations must be solved M times. Thus, each iteration of the gradient method has complexity:

$$C_{\text{TD, gradient}} = \mathcal{O}(M^3 \cdot \kappa N_c \cdot \log_2(\kappa N_c)) \quad (20)$$

To obtain more consistent timing, we used a fixed number ($=20$) of iterations for the gradient descent solver, instead of a stopping criterion based on error.

Finally, the frequency-domain JLS filter synthesis inverts a size M matrix for each of roughly κN_c frequency points, so complexity scales as:

$$C_{\text{FreqDomain}} = \mathcal{O}(M^3 \kappa N_c) \quad (21)$$

The complexities for the MMF constructions are much less than those for JLS, because the filter for each waveform of the CPI is synthesized independently of the others. Thus the scaling with number of waveforms is $\mathcal{O}(M)$ for MMF instead of $\mathcal{O}(M^3)$ for JLS. The scalings with κN_c are the same as for JLS filters.




Fig. 6. Scaling of complexity for three JLS synthesis methods versus number of waveforms M , for fixed filter length $\kappa N_c = 1800$.

V. CONCLUSION

We have shown that MMF and JLS pulse-compression filters, which have previously been defined via a time-domain system of equations, can also be synthesized via a more efficient frequency-domain representation. This is an important step toward just-in-time computation of pulse-compression filters for radars employing pulse-to-pulse agility.

REFERENCES

- [1] Richards, M., Scheer, J., and Holm, W., Eds., *Principles of Modern Radar*, vol. 1, SciTech Publishing, 2010.
- [2] Richards, M., *Fundamentals of Radar Sig. Proc.*, McGraw-Hill, 2005.
- [3] Levanon, N. and Mozeson, E., *Radar Signals*, John Wiley & Sons, 2004.
- [4] Scholnik, D., "Range-ambiguous clutter suppression with pulse-diverse waveforms," in *IEEE Radar Conf.*, May 2011, pp. 336–341.
- [5] Jakabosky, J., Blunt, S., and Himed, B., "Spectral-shape optimized FM noise radar for pulse agility," in *IEEE Radar Conf.*, June 2016, pp. 1–6.
- [6] Blunt, S., Cook, M., and Stiles, J., "Embedding information into radar emissions via waveform implementation," in *IEEE Intl. Waveform Div. Design Conf.*, August 2010, pp. 195–199.
- [7] Davis, M. and Pillai, S. U., "Waveform diversity for ultra-wide band surveillance radars," *IET Radar Sonar Navig.*, 8(9): 1226–1233, 2014.
- [8] Mokole, E. and Wicks, M., "A short history of waveform diversity," in Wicks, M., Mokole, E., Blunt, S., Schneible, R., and Amuso, V. [Eds.], *Principles of Waveform Diversity and Design*, SciTech Publishing, 2010.
- [9] Griffiths, H., Blunt, S., Cohen, L., and Savy, L., "Challenge problems in spectrum engineering and waveform diversity," in *IEEE Radar Conf.*, May 2013, pp. 1–5.
- [10] Blunt, S. and Mokole, E., "Overview of radar waveform diversity," *IEEE Aerosp Electron Sys Mag*, 31(11): 2–42, 2016.
- [11] O'Connor, A., Kantor, J., and Jakabosky, J., "Joint equalization filters that mitigate waveform-diversity modulation of clutter," in *IEEE Radar Conf.*, June 2016, pp. 1–6.
- [12] O'Connor, A., Kantor, J., and Jakabosky, J., "Filters that mitigate waveform modulation of radar clutter," *IET Radar Sonar Navig.*, 11(8): 1188–1195, 2017.
- [13] Ackroyd, M. and Ghani, F., "Optimum mismatched filters for sidelobe suppression," *IEEE Trans. Aerosp. Electron. Sys.*, 9(2): 214–218, 1973.
- [14] Davis, R., Fante, R., and Perry, R., "Phase-coded waveforms for radar," *IEEE Trans. Aerosp. Electron. Sys.*, 43(1):401–408, 2007.
- [15] Stoica, P., Li, J., and Xue, M., "Transmit codes and receive filters for radar," in *IEEE Sig. Proc. Mag.*, 25(6): 94–109, 2008.
- [16] Oppenheim, A., Schaffer, R., and Buck, J. *Discrete-Time Signal Processing*, 2nd Ed., Prentice Hall, 1998.
- [17] Adali T. and Li, H., "Complex-valued adaptive signal processing," in T. Adali and S. Haykin, [Eds.], *Adaptive Signal Processing, Next Generation Solutions*, John Wiley & Sons, 2010.
- [18] Blake, W., "Multi-dimensional mismatch filter design optimization for radar waveforms," in *IEEE Radar Conf.*, June 2016, pp. 1–4.